Distributed Constraint Problems for Utilitarian Agents with Privacy Concerns, Recast as POMDPs

Julien Savaux, Julien Vion, Sylvain Piechowiak, René Mandiau, *LAMIH UMR CNRS 8201, University of Valenciennes, France*
Toshihiro Matsui, *Nagoya Institute of Technlogy, Japan*
Katsutoshi Hirayama, *Kobe University, Japan*
Makoto Yokoo, *Kyushu University, Japan*
Shakre Elmane, Marius Silaghi, *Florida Institute of Technology, USA*

arXiv:1703.06939v1 [cs.AI] 20 Mar 2017

Abstract

Privacy has traditionally been a major motivation for distributed problem solving. Distributed Constraint Satisfaction Problem (*DisCSP*) as well as Distributed Constraint Optimization Problem (*DCOP*) are fundamental models used to solve various families of distributed problems. Even though several approaches have been proposed to quantify and preserve privacy in such problems, none of them is exempt from limitations. Here we approach the problem by assuming that computation is performed among utilitarian agents. We introduce a utilitarian approach where the utility of each state is estimated as the difference between the reward for reaching an agreement on assignments of shared variables and the cost of privacy loss. We investigate extensions to solvers where agents integrate the utility function to guide their search and decide which action to perform, defining thereby their policy. We show that these extended solvers succeed in significantly reducing privacy loss without significant degradation of the solution quality.

I. Introduction

Privacy is an important problem in a lot of distributed applications, therefore, the reward for solving the problem should be considered, but also the cost of privacy loss during the process [22]. For example, when users exchange information on social networks [32], they reveal (often unconsciously) personal data (*e.g.*, age, address, date of birth). Related studies were also performed in the domain of Ambient Intelligence which is concerned with the distributed management of confidential information between different components (*e.g.*, camera, computer, PDA), to allow or prevent the sharing of resource information [47].

In distributed scheduling problems, confidentiality also happen when information is exchanged between agents/participants (for example, a participant does not want to reveal his/her unavailability for a time slot because the explanation concerns his/her private life, and it should not be necessarily discussed publicly). Indeed, we know that the assignment of time slots can be difficult if participants do not want to reveal their constraints [18], [10]. Such coordinated decisions are in conflict with the need to keep constraints private [16]. We consider that these coordinated decisions may be defined by a set of different constraints distributed among different agents.

In Distributed Constraint Satisfaction Problems (*DisCSPs*) and Distributed Constraint Optimization Problems (*DCOPs*), agents have to assign values to variables while respecting given constraints. To find such assignments, agents exchange messages until a solution is found or until an agent detects that there is none. Thus, agents have to reveal information during search, causing privacy to be a major concern in *DisCSPs* [62], [26]. If an agent is concerned about its privacy, then it can associate a cost to the revelation of each information in its local problem. This cost may be embedded into utility driven reasoning.

A common assumption is that utility-based agents associate each state with a utility value [48]. As such, each communication action's utility is evaluated as the difference between initial and final utilities. Indeed, if an agent is concerned about its privacy, then it can associate a cost to the revelation of each information in its local problem. Since they are interested in solving the problem, they must also be able to quantify the reward they draw from finding the solution. Here we approach the problem by assuming that privacy is a utility that can be aggregated with the reward for solving this problem. We evaluate how much privacy is lost by the agents during the problem solving process, by the sum of the utility lost for each information that was revealed. For example, the existence and availability of a value from the domain of a variable is the kind of information that the agents want to keep private. The cost of a constraint for a solution is another example of information that agents would like to keep private.

While sometimes possibilistic reasoning is used to guide search, agents were usually assumed to participate in the search process until an agreement is found between all different agents (*i.e.*, a solution). We investigate the case where an agent may modify its search process to optimize utility. Two extensions are introduced, Utilitarian Distributed Constraint Satisfaction Problem (*UDisCSP*) and Utilitarian Distributed Constraint Optimization Problem (*UDCOP*), addressing *DisCSPs* and *DCOPs*, respectively. These extensions exploit the rewards of agreements and costs representing privacy loss as guidance for the utility-based agents, where the utility of each state is estimated as the difference between the expected rewards for agreements, and the expected cost representing privacy loss. In this work, Distributed Constrained Problem (*DisCP*) will refer to both *DisCSP* and *DCOP*. Similarly, Utilitarian Distributed Constrained Problems (*UDisCPs*) will refer to both *UDisCSPs* and *UDCOPs*.

The paper is organized as follows. Section II presents existing research concerning solving algorithms and approaches to privacy for *DisCPs*. Further, Section III describes the concepts and solvers involved in *UDisCP* to preserve privacy. Section IV discusses theoretical implications. Section V reports our experimental results and evaluates privacy loss on distributed meeting scheduling (DMS) problems. Section VI presents our conclusions and our directions for future research.

II. Background

We present distributed constrained problems (II-A), and existing approaches for privacy as well as their limits (II-B).

A. Distributed Constrained Problem

Distributed Constraint Satisfaction Problems (*DisCSPs*) and Distributed Constraint Optimization Problems (*DCOPs*) have been extensively studied as a fundamental way of modelling constrained problems in multi-agent systems, and will be defined in the following, as well as existing solvers.

1) Definitions: Let us first remind the definitions of the Distributed Constraint Satisfaction Problem (*DisCSP*) and of the Distributed Constraint Optimization Problem (*DCOP*) [62].

Definition 1. *A Distributed Constraint Optimization Problem (DCOP) is formally defined as a tuple $\langle \mathscr{A}, \mathscr{X}, \mathscr{D}, \mathscr{C} \rangle$ where:*

- *$\mathscr{A} = \{A_1, A_2, \ldots, A_m\}$ is a finite set of m agents.*
- *$\mathscr{X} = \{x_1, x_2, \ldots, x_n\}$ is a finite set of n variables. Each agent A_{self} encapsulates variables denoted $\mathscr{X}(A_{self})$ (with $\mathscr{X}(A_{self}) \subseteq \mathscr{X}$).*
- *$\mathscr{D} = \langle \mathscr{D}_1, \mathscr{D}_2, \ldots, \mathscr{D}_n \rangle$ is a set of n domains. $\mathscr{D}_{self}$ is the set of possible values for $\mathscr{X}(A_{self})$.*
- *$\mathscr{C} = \{C_1, C_2, \ldots, C_e\}$ is a finite set of e valued constraints. Each constraint C_i involves some variables $\mathscr{X}(C_i) \subseteq \mathscr{X}$ defining a cost (positive value) for assignments. We note that $\mathscr{C}_{self} = \{C_i \in \mathscr{C} | \mathscr{X}(C_i) \cap \mathscr{X}(A_{self}) \neq \emptyset\}$.*

The objective is to find an assignment for each variable that minimizes the total cost.

Definition 2. *A Distributed Constraint Satisfaction Problem (DisCSP) is defined as a DCOP where the constraints are predicates, each one defining for sets of assignments a cost of ∞ (constraint violation). The problem is to find an assignment to variables that does not violate any constraint.*

Definition 3. *Distributed Constrained Problem (DisCP) is any problem modelled with DisCSP or one of its extensions.*

In the following, we study a particular case of problems, namely mono-variable problems where $m = n$ (*i.e.*, a variable per agent), $\mathscr{X}(A_{self})$ containing a single variable, x_{self}. Also, in the following frameworks and algorithms, x_{self}, $\mathscr{D}_{self}$, and $\mathscr{C}_{self}$ are generalized referring to the projection of the part of different elements from the problem known only by the current agent (called $self$).

The problem that each agent has to solve in *DisCP* is a Stochastic Constraint Optimization Problem (Stochastic COP) or a Stochastic CSP [60], which are generalizations of Stochastic SAT, SSAT [34], namely where the local problem has to be solved considering its impact on external variables and constraints whose values are not yet known and that are not under the control of this agent, but that of other agents (commonly, lower priority), values of which the agent may know a probabilistic profile.

2) Existing Solvers: We now introduce some existing solvers for *DisCP* with which we exemplify extensions based on our utilitarian approach. We consider two well-known solvers for *DisCSP* (*SyncBT* and *ABT*), and three solvers for *DCOPs* (called *ADOPT*, *DSA* and *DBO*):

a) Synchronous Backtracking (SyncBT): is the baseline algorithm for *DisCSPs* [64], [68]. *SyncBT* is a simple distribution of the standard backtracking algorithm. Agents consecutively send a satisfying assignment to their variable to the next agent. If an agent is unable to find an instantiation compatible with the current partial assignment it has received, it asks the previous one to change its assignment. The process repeats until a complete solution is built or until the whole search space is explored.

b) Asynchronous Backtracking (ABT): is a common alternative solver for *DisCSPs* that allows agents to run concurrently and asynchronously [64]. Each agent finds an assignment to its variable and communicates it to the other connected agents having constraints involving this variable. Then agents wait for incoming messages. The assignments received through `ok?` form a context called `agentView`. If an agent's assignment is inconsistent with its `agentView`, it is changed and communicated to the other agents. A subset of an `agentView` preventing an agent from finding an assignment that does not violate any of its constraints is called a nogood. If an agent infers a nogood from its constraints and its `agentView`, it asks the lowest priority agent involved in the nogood to change its assignment through a `nogood` message.

c) Asynchronous Distributed Optimization (ADOPT): guarantees to find the optimal solution and only needs polynomial space [42], [61], [56]. *ADOPT* organizes agents into a depth first search tree in which constraints are only allowed between a variable and any of its acquaintances (parents and descendants in the tree). When the estimated cost of agents' assignment is higher than a given threshold, the agent switches its value assignment to the value with smallest estimated cost. When the upper and lower bounds meet at the root agent, a globally optimal solution has been found and the process finishes. Note that other approaches based on a graph re-arrangement have been explored, for example a cluster exploitation like Asynchronous Partial Overlay [39], [38].

d) Distributed Stochastic Algorithm (DSA): makes agents start their process by randomly selecting a value [67]. Agents then enter an iteration where they send their last assigned value to their neighbours and collect any new values from them. The next candidate value is chosen based on the values received from other agents and on maximizing a given utility function. This algorithm is incomplete and does not guarantee optimality, but it is frequently efficient for finding solutions close to optimum.

e) Distributed Breakout (DBO): is an iterative improvement solver for *DCOPs* [63], [28]. The evaluation of a given solution is the summation of the weights for all its violated constraints. An assignment is then changed to decrease the solution value. If the evaluation of the solution cannot be decreased by changing a value, the current state may be a local minimum. In this situation, *DBO* increases the weights of constraint violation pairs in the current state so that the evaluation of this current state becomes higher than the neighbouring states. Thus, the algorithm can escape from the local minimum.

B. Privacy

We present generalities about privacy, in particular about our definition and typology. This typology will be used to compare existing approaches to privacy in *DisCP*.

1) Definition and typology: Privacy is the concern of agents to not reveal their personal information. In this work, we define privacy as follows:

Definition 4. *Privacy is the utility that agents benefit from conserving the secrecy of their personal information.*

Contrary to the standard rewards in *DisCSPs*, privacy costs are proper to each individual agent. Therefore, the computation is now performed by utility-based and self-interested agents, whose decisions aims at maximizing a utility function. The objective is then to define a policy associating an expected utility maximizing action (communication act or computation) to each state, where the state includes the belief about the global state). In existing works, several approaches have been developed to deal with privacy in *DCOPs*. Some cryptographic approaches offer certain end-to-end security guarantees by integrating the entire solving process in one primitive for *DisCSPs* or for *DCOPs*, the highest level of such privacy guarantees being achievable only for problems with a single variable [54]. Other cryptographic approaches are hybrids interlacing cryptographic and artificial intelligence steps [57], [21], [24]. While ensuring privacy in each primitive [29], cryptographic techniques are usually slower, and sometimes require the use of external servers or computationally intensive secure function evaluation techniques that may not always be available or justifiable for their benefits, making them impractical [22], or lacking clear global security guarantees. A couple of such approaches with which we compare in more detail are:

a) Distributed Pseudo-tree Optimization Procedure with Secret Sharing (SS-DPOP): modifies the standard *DPOP* algorithm [46] to protect leaves in the depth first search tree, where agents sharing constraints are on the same branch [21]. *SS-DPOP* uses secret sharing [52] to aggregate the results of a single solution, without revealing the individual valuations this solution consists in. The aggregated values are then passed to the bottom agent, who aggregates this information with its own valuations and sends the aggregate up the chain. [33] has also extended *DPOP* to preserve different types of privacy using secure multi-party computation.

b) Privacy-Preserving Synchronous Branch and Bound (P-SyncBB): is a cryptographic version of *SyncBB* [27], [25] for solving *DCOPs* while enforcing a stronger degree of constraint privacy [24]. *P-SyncBB* computes the costs of CPAs (current partial assignments) and compares them to the current upper bound, using secure multi-party protocols. Some protocols were proposed in [24] that can solve problems without resorting to costly transfer sub-protocols, and compare the cost of a CPA shared between two agents to the upper bound held by only one of them. Some variations on the standard *SyncBB* include *NCBB* and *AFB* [20], [24].

We choose to deal with privacy by embedding it into agents' decision-making. Other approaches use different metrics and frameworks to quantify privacy loss. According to [23], agents privacy may concern the four following aspects:

> Domain privacy: Agents want to keep the domain of their variable private. The common benchmarks and some algorithms (like *DPOP*, and common cryptographic techniques) assume that all the domains are public, which leads to a complete loss of domain privacy. In the original *DisCSP* approach a form of privacy of domains is implicit (see *ABT*), while being formally required in its *PKC* extension [6].
>
> Constraint privacy: Agents want to keep the information related to their constraints private [53]. If variables involved in constraints are considered to belong to only one agent, we can distinguish the revelation of information to agents that participate in the constraint (internal constraint privacy) and the one to other agents (external constraint privacy). While common problems with domain privacy can be straightforwardly modelled as problems with constraint privacy, as discussed later, there theoretically exists a kind of domain privacy which cannot be modelled with constraint privacy.
>
> Assignment privacy: Agents want to keep the assigned values to their variables private. The revelation of assigned values concerns the assignment of the final solution, as well as the ones proposed during search.
>
> Algorithmic privacy: Even though it is commonly assumed that all agents run the same algorithm during the solving, agents may modify the value of some parameters guiding the search process for some personal benefit (*e.g.*, the likelihood of updating its value). This can be achieved by keeping the message structure and contracts of certain existing *DisCSP* solvers to be used as communication *protocols* rather than *algorithms*, as introduced in [55], where protocols obtained in such ways are compared with respect to the flexibility offered for agents to hide their secrets.

2) Existing approaches to privacy in DisCP: We now introduce two examples modelled with *DisCP* frameworks, and show how existing approaches to privacy deal with these problems and their limits.

a) Distributed Constraint Satisfaction Problem:

Example 1. *Suppose a meeting scheduling problem between a professor (called A_1) and two students (A_2 and A_3). They all consider to agree on a time slot to meet on a given day, having to choose between $8\ am$, $10\ am$ and $2\ pm$. Professor A_1 is unavailable at $2\ pm$, A_2 is unavailable at $10\ am$, and A_3 is unavailable at $8\ am$.*

There can exist various reasons for privacy. For example, A_2 does not want to reveal the fact that it is busy at $10\ am$. The value that A_2 associates with not revealing the $10\ am$ unavailability is the salary from a second job ($\$2{,}000$). The utility of finding an agreement for each student is the stipend for their studies ($\$5{,}000$). For A_1, the utility is a fraction of the value of its project ($\$4{,}000$). This is an example of privacy for absent values or constraint tuples.

Further A_3 had recently boasted to A_2 that at $8\ am$ it interviews for a job, and it would rather pay $\$1{,}000$ than to reveal that it is not. This is an example of privacy for feasible values of constraint tuples.

Similarly, participants associate a cost to the revelation of each availability and unavailability. Thus, scaling numbers by 1000 for simplicity, corresponding agents associate a cost of 1 to the revelation of their availability at $8\ am$, a cost of 2 to the one at $10\ am$, and a cost of 4 to the one at $2\ pm$. The reward from finding a solution is 4 for A_1 and 5 for both A_2 and A_3.

For simplicity, in the next sections, we will refer to the possible values by their identifier: 1, 2, and 3 (corresponding to $8\ am$, $10\ am$ and $2\ pm$ respectively). As this problem states allowed or forbidden values, it is represented by a *DisCSP* as follows:

- $\mathscr{A} = \{A_1, A_2, A_3\}$
- $\mathscr{X} = \{x_1, x_2, x_3\}$
- $\mathscr{D} = \{\{1,2,3\}, \{1,2,3\}, \{1,2,3\}\}$
- $\mathscr{C} = \{\{\neg(x_1 = x_2 = x_3)\}, \{(x_1 \neq 3)\}, \{(x_2 \neq 2)\}, \{(x_3 \neq 1)\}\}$

As it can be observed, *DisCSPs* cannot model the information concerning privacy. Now we will show how existing extensions model them.

b) Distributed Private Constraint Satisfaction Problem (DisPrivCSP): models the privacy loss for individual revelations [18], [55]. It also lets agents abandon the search process when the incremental privacy loss overcomes the expected gains from finding a solution. Each agent pays a cost if the feasibility of each solution is determined by other agents. The reward for solving the problem is given as a constant. Those concepts were so far used for evaluating qualitatively existing algorithms, but were not integrated as heuristics in the search process. Privacy and the usual optimization criteria of Distributed Constraint Optimization Problems are merged into a unique criterion [15]. The additional parameters are a set of privacy coefficients and a set of rewards.

This framework successfully models all the information described in the initial problem and also measures the privacy loss for each agent. However, it was not yet investigated what is the impact of the interruptions when privacy loss exceeds the reward threshold, its relation to utility, or how agents could use this information to modify their behaviour during the search process to preserve more privacy.

c) Valuation of Possible States (VPS): measures privacy loss by the extent to which the possible states of other agents are reduced [58], [36], [35], [22]. Privacy is interpreted as a valuation on the other agents' estimates about the possible states that one lives in. During the search process, agents propose their values in an order of decreasing preference. At the end of the search process, the difference between the presupposed order of preferences and the real one observed during search determines the privacy loss: the greater the difference, the more privacy has been lost.

However in our sample problem, agents initially know nothing about others agents but the variable they share a constraint with and cannot suppose an order of preference. Agents have no information about others agents privacy requirements. Thus, agents do not expect to receive any value proposal more than another. In this direction one needs to extend *VPS* to be able to also model the kind of privacy introduced in this example.

d) Partially Known Constraints (PKC): uses entropy, as defined in information theory, to quantify privacy loss [6]. In this method, two variables owned by two different agents may share a constraint. However, not all the forbidden couples of values involved in a constraint are known by both agents. Each agent only knows a subset of the constraints. During the search process, assignment privacy is leaked through `ok?` and `nogood` messages, like in standard algorithms. This problem is solved by not sending the value that is assigned to a variable in a `ok?` message, but the set of values compatible with this assignment. For `nogood` messages, rather than sending the current assignments, an identifier is used to specify the state of the resolution and to check if some assignments are obsolete or not.

PKC assumes that agents only know their own individual unavailabilities [6]. Only the junction of information known by all agents can rebuild the whole problem. However, while *PKC* let agents preserve privacy of unary constraints, it does not consider the cost of revelation of assignments.

e) Distributed Constraint Optimization Problem:

Example 2. *Suppose a problem concerning scheduling a meeting between three participants. They all consider to agree on a place to meet on a same slot time, to choose between London, Madrid and Rome. For simplicity, we will refer to these possible values by their identifiers:* 1*,* 2 *and* 3 *(London, Madrid and Rome respectively).*

A_1 *lives in Paris, and it will cost it* \$70*,* \$230 *and* \$270 *to attend the meeting in London, Madrid and Rome respectively.* A_2 *lives in Berlin, and it will cost it* \$120*,* \$400 *and* \$190 *to attend the meeting in London, Madrid and Rome respectively.* A_3 *lives in Brussels, and it will cost it* \$40*,* \$280 *and* \$230 *to attend the meeting in London, Madrid and Rome respectively. The objective is to find the meeting location that minimizes the total cost students have to pay in order to attend.*

The privacy costs for revealing its cost for locations 1*,* 2 *and* 3 *for* A_1 *are* $\$80, \$20, \$40$*. The privacy cost for locations* 1*,* 2 *and* 3 *are* $\$100, \$30, \$10$ *for* A_2 *and* $\$80, \$30, \$10$ *for* A_3*. There exist various reasons for privacy. For example, students may want to keep their cost for each location private, since it can be used to infer their initial location, and they would pay an additional (privacy) price rather than revealing the said travel cost. For example,* A_1 *associates* \$80 *privacy cost to the revelation of the travel cost of* \$70 *for meeting in London.*

This example may be defined by a *DCOP*:

- $\mathscr{A} = \{A_1, A_2, A_3\}$
- $\mathscr{X} = \{x_1, x_2, x_3\}$
- $\mathscr{D} = \{\{1,2,3\}, \{1,2,3\}, \{1,2,3\}\}$
- $\mathscr{C} = \{\{(x_1 = 1), 70\}, \{(x_1 = 2), 230\}, \{(x_1 = 3), 270\},$
  $\{(x_2 = 1), 120\}, \{(x_2 = 2), 400\}, \{(x_2 = 3), 190\},$
  $\{(x_3 = 1), 40\}, \{(x_3 = 2), 280\}, \{(x_3 = 3), 230\},$
  $\{\neg(x_1 = x_2 = x_3), \infty\}\}$

The notation $(x = a)$ is a predicate p stating that variable x is assigned to value a. Each constraint in $\mathscr{C}$ is described with the notation $\{p, v_{i,p}\}$, and states that if predicate p holds then a cost $v_{i,p}$ is paid by the Agent A_i enforcing the constraint.

One could attempt to model the privacy requirements by aggregating the solution quality, related to obtain the reward for a solution (called $SolutionCost$) and for keeping the privacy (called $PrivacyCost$) into a unique cost. However, this is not possible. Indeed, in a *DCOP*, agents explore the search space to find a better solution, and only pay the corresponding solution cost when the search is over and the solution is accepted. This means that the solution cost decreases with time. However, privacy costs are cumulative and are paid during the search process itself (at each time, a solution is proposed), no matter what solution is accepted at the end of the computation. This means that the total privacy loss increases with time. Aggregating the solution costs and privacy costs or using a multi-criteria *DCOP* would not consider the privacy cost of the solutions that are proposed but not kept as final. Also, a given solution may imply different privacy losses depending on the algorithm used to reach it.

None of the previous techniques consider all aspects of privacy, and while they preserve some privacy, related algorithms require resources or properties from the *DisCSP* or *DCOP* that may not be possible due to requirements dictated by the real world problem. We therefore propose to deal with privacy directly from the agents' decision making perspective, using a utility-based approach.

III. Utilitarian Distributed Constrained Problem

This section defines our utility-based framework (III-A), and then describes how it models previously presented problems, with different types of privacy requirements (III-B).

A. Definitions

While some previously described frameworks do model the details of our example regarding privacy, it has until now been an open question as to how they can be dynamically used by algorithms in the solution search process. It can be noticed that the rewards and costs in our problem are similar to the utilities commonly used by planning algorithms [31]. Thus, we propose to define a framework that specifies the elements of the corresponding family of planning problems. To do so, we ground the theory of our interpretation of privacy in the well-principled theory of utility-based agents [51].

Definition 5. *A utility-based agent is characterized by its ability to associate a value to each state of the problem, representing its contentment to be in this state.*

For example, the *state* of Agent A_i can include the subset of $\mathscr{D}_i$ that it has revealed and the reward associated to the solution under consideration. The problem is to define a policy for each agent such that their utility is maximized. A policy is a function that associates each state with an action that should be performed in it [48]. We evaluate the utility of a state as follows:

$$Utility = \sum_i reward_i - \sum_j cost_j \text{ where } reward_i, cost_j \in \mathbb{R}^+ \quad (1)$$

Possible actions are communications dictated by the used solver, and each possible solution proposal, acceptation or rejection, is associated with its corresponding reward (for increasing the probability of problem solving) and cost (for revealing information).

Thus, unlike for *DisCPs*, the solution of a *UDisCP* does not necessarily include an agreement, as pursuing the search may imply a decrease of utility. As privacy is lost during search, a given set of assignments can have different utilities, depending on the information exchanged before. We define Utilitarian Distributed Constrained Problem (*UDisCP*):

Definition 6. *A Utilitarian Distributed Constrained Problem (UDisCP) is formally defined as a tuple* $\langle \mathscr{A}, \mathscr{X}, \mathscr{D}, \mathscr{C}, \mathscr{U}, \mathscr{R} \rangle$*, i.e., a DisCP with:*

- $\mathscr{U} = \langle \mathscr{U}_d, \mathscr{U}_a, \mathscr{U}_g, \mathscr{U}_c \rangle$ *is a quadruplet where:*
  - $\mathscr{U}_d$ *is a matrix of domain privacy costs, where* $u_{d(i,j)}$ *is the cost for* A_i *to reveal whether* $j \in \mathscr{D}_i$*.*
  - $\mathscr{U}_a$ *is a matrix of assignment privacy costs, where* $u_{a(i,j)}$ *is the cost for* A_i *to reveal a local solution including the assignment* $x_i = j$*.*
  - $\mathscr{U}_g$ *is a matrix of assignment privacy costs, where* $u_{g(i,j)}$ *is the cost for* A_i *to reveal the existence of a global solution including the assignment* $x_i = j$*.*
  - $\mathscr{U}_c$ *is a vector of constraint privacy costs, where* $u_{c(i,j)}$ *is the cost for* A_i *to reveal the weight of constraint* C_j*.*
- $\mathscr{R} = \{r_1, \ldots, r_n\}$ *is a vector of rewards, where* r_i *is the reward that Agent* A_i *receives if an agreement is found.*

Note that is the case where a domain can be represented on a computer as a type of a parameter to a predefined software function modelling constraints (*e.g.*, an integer), then $\mathscr{U}_d$ can be embedded in $\mathscr{U}_c$, as a unary constraint on the domain. We note also the notations *UDisCSP* for Utilitarian *DisCSP* and *UDCOP* for Utilitarian *DCOP*.

B. Description on Problems with Privacy

In this section, we present how *UDisCP* deals with the question of privacy on the two previous examples, and we solve them using our utilitarian extensions of existing solvers. For clarity, unchanged parts of pseudo-codes are shown in gray colour (while our extensions are in black colour). These parts may include variables or procedures that are used in other parts of the solvers but that are not used (and therefore not detailed) in this work.

1) Privacy of Domains: We propose to revise our previous example and present the extensions for two standard solvers (*SyncBT* and *ABT*).

Example 3. *Recall* $\mathscr{U}_d$ *represents the cost for each agent to reveal possible values.* $\mathscr{R}$ *is the reward that each agent gets when a solution is found, motivating them to initiate the solving.*

DisCSP introduced in Example 1 is extended to UDisCSP by specifying the additional parameters $\mathscr{U}_d$ *and* $\mathscr{R}$ *defined as follows. For example,* $u_{d(1,3)}$ *refers to the cost for the first agent to reveal the third value in its domain, namely* 4*.*

$$\mathscr{U}_d = \begin{bmatrix} 1 & 2 & 4 \\ 1 & 2 & 4 \\ 1 & 2 & 4 \end{bmatrix} \begin{matrix} A_1 \\ A_2 \\ A_3 \end{matrix} \qquad \mathscr{R} = \{4, 5, 5\}$$

Now we discuss how the standard *ABT* and *SyncBT* algorithms are adjusted to *UDisCSPs*. After each state change, each agent computes the estimated utility of the state reached by each possible action, and selects randomly one of the available actions leading to the state with the maximum expected utility.

SyncBTU and *ABTU* are obtained by similar modifications of *SyncBT* and *ABT* [64], [62], [68], respectively. In our extended algorithms, agents compute the frequency of rejection of their solution proposal to estimate the expected utilities. This frequency can be re-evaluated at any moment based on data recorded during previous runs on problems of similar tightness (*i.e.*, having the same proportion of forbidden instantiations). Learning from previous experience has been extensively studied [1], [2]. The learning can be off-line or on-line. For off-line learning, agents calculate the number of messages `ok?` and `nogood` sent during previous executions, called $count$. They also count how many messages previously sent lead to the termination of the algorithm, in the variable $agreementCount$. The frequency with which a solution leads to the termination of the algorithm, is called $agreementProb$ (Equation 2). For on-line learning, one can update the variables $count$, $agreementCount$ and $agreementProb$ dynamically whenever the corresponding events happen. When previous experiments are not available, the value of $agreementProb$ is set to $1/2$ by default (this value is always between 0 and 1).

$$agreementProb = \frac{agreementCount}{count} \tag{2}$$

When `ok?` messages are sent, the agent has the choice of which assignment to propose. When a `nogood` message is scheduled to be sent, agents also have choices of how to express them. Before each `ok?` or `nogood` message, the agents check which available action leads to the highest expected utility. If the highest expected utility is lower than the current one, the agent announces failure. The result is used to decide between proposing assignments, a nogood, or declaring failure.

a) Synchronous Backtracking with Utility (SyncBTU): *SyncBTU* is obtained by restricting the set of actions to the standard communicative acts of *SyncBT*, namely `ok?` and `nogood` messages. The procedures of a solver like *SyncBT* define a policy, since they only identify a set of actions (inferences and communications) to be performed in each state. A state of an agent in *SyncBT* is defined by `agentView` and a current assignment of the local variable. The local inferences in *SyncBTU* are obtained from the ones of *SyncBT* by an extension exploiting the utility information available. The criteria in this research was not to guarantee an optimal policy but to use utility with a minimal change to the original behaviour of *SyncBT* reinterpreted as a policy. In *SyncBTU*, the state is extended to also contain a history of revelations of one's values defining an accumulated privacy loss, and a probability to reach an agreement with each action. Since [64] does not provide pseudo-code for *SyncBT*, we modify the pseudo-code presented in [68] for *SyncBTU*: `assignCPA` (before Line 7), and `backtrack` (before Line 6). As *SyncBTU* is close from *ABTU*, we do not detail the pseudo-code for this algorithm in this report.

b) Asynchronous Backtracking with Utility (ABTU): Similar modifications are applied to *ABT* to obtain *ABTU*: communications of *ABTU* are composed of `ok?`, `addlink` and `nogood`. The state and local inferences of *ABTU* are similar to *SyncBTU*, while also containing the set of nogoods.

Algorithm 1: `checkAgentView_ABTU`

Input: $\mathscr{D}_{self}$, `agentView`, $agreementProb$, r_{self}
1 **when** $agentView$ and $currentValue$ are inconsistent **do**
2 **if** *no value in* $\mathscr{D}_{self}$ *is consistent with* `agentView` **then**
3 `backtrack`;
4 **else**
5 select $d \in \mathscr{D}_{self}$ where `agentView` and d are consistent;
6 $currentValue \leftarrow d$;
7 **if** (`estimateCostDisCSP` $(agreementProb, \mathscr{D}_{self}, 1) \geqslant r_{self}$) **then**
8 `interruptSolving`;
9 **else**
10 send (**ok?**,(x_{self};d)) to outgoing links

To calculate the estimated utility of pursuing an agreement (revealing an alternative assignment), the agent considers all different possible scenarios of the subsets of values that might have to be revealed in the future based on possible rejections received, together with their probability (Algorithm 2). This algorithm assumes as parameters: (i) $agreementProb$ (Equation 2), (ii) an ordered set of possible values $\mathscr{D}'_{self}$ for a scenario (by default, the order proposed in $\mathscr{D}_{self}$), and (iii) the probability to select a value from $\mathscr{D}'_{self}$, initially 1), called $probD$. Note that $\mathscr{D}'_{self}[j]$ refers to value at index j of $\mathscr{D}'_{self}$.

Algorithm 2: `estimateCostDisCSP`

Input: $agreementProb$, $\mathscr{D}'_{self}$, $probD$
Output: $estimatedCost$
1 $valueId = j \mid (\mathscr{D}_{self}[j] = \mathscr{D}'_{self}[1])$;
2 **if** ($|\mathscr{D}'_{self}| = 1$) **then**
3 **return** $(\sum_{j=1}^{j=valueId} u_{d(self,j)}) \times probD$;
4 **else**
5 $v \leftarrow \mathscr{D}'_{self}[1]$;
6 $costRound \leftarrow$ `estimateCostDisCSP` ($agreementProb$, $\{v\}$, $agreementProb \times probD$);
7 $costTemp \leftarrow$ `estimateCostDisCSP` ($agreementProb$, $\mathscr{D}'_{self} \setminus \{v\}$, $(1 - agreementProb) \times probD$);
8 $estimatedCost \leftarrow costRound + costTemp$;
9 **return** $estimatedCost$;

The algorithm then recursively computes the utility of the next possible states, and whether the revelation of the current value v leads to the termination of the algorithm, values which are stored in variables $costRound$ and $costTemp$. The algorithm returns the estimated cost (called $estimatedCost$) of privacy loss for the future possible states currently.

Example 4. *Continuing with Example 1 (whose a solving is illustrated by Figure 1), at the beginning of the solving, Agent* A_1 *has to decide for a first action to perform. We suppose the* $agreementProb$ *learned from previous solvings is* 0.5. *To decide whether it should propose an available value or not, it calculates the corresponding* $estimatedCost$ *by calling Algorithm 2 with parameters: the learned* $agreementProb = 0.5$*, the set of possible solutions (*$\mathscr{D}'_1 = \{1, 2, 3\}$*) and* $probD = 1$.

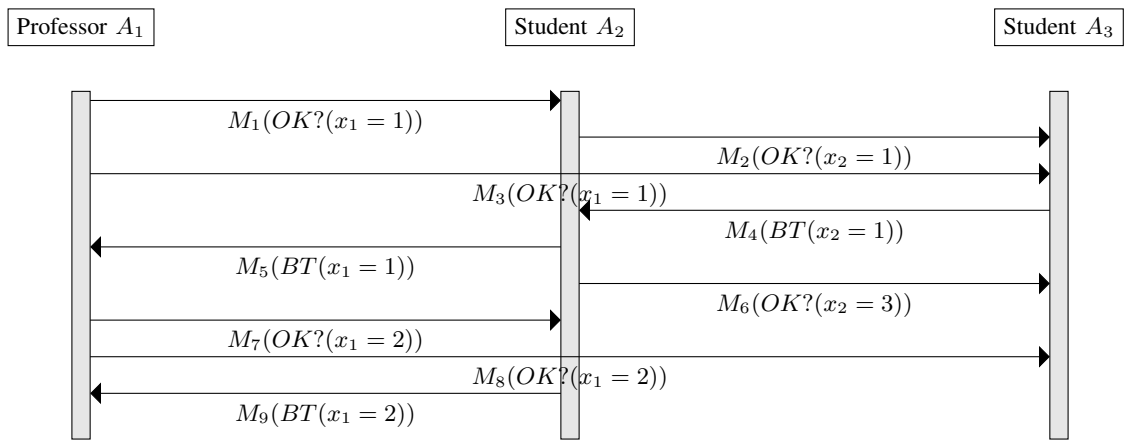

Fig. 1. Interactions between agents during *ABT*

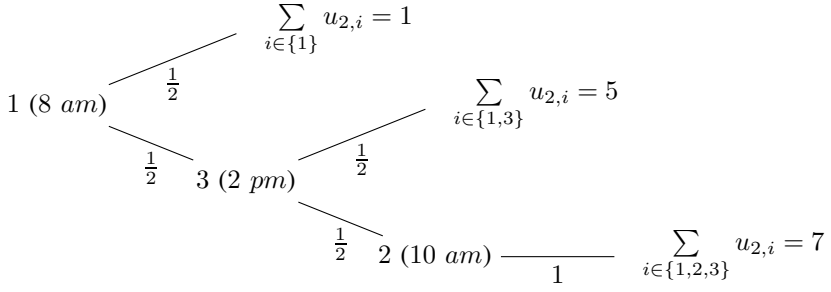

Fig. 2. Cost estimation for all scenarios by A_2 during *ABTU*

For each possible value, this algorithm recursively sums the cost for the two scenarios corresponding to whether the action leads immediately to termination, or not. Given privacy costs, the availability of three possible subsets of $\mathscr{D}_1'$ may be revealed in this problem: $\{1\}, \{1,2\}$, *and* $\{1,2,3\}$. *Each set of size S consists of S first elements of the list solution based on this initial order.*

The $estimatedCost$ *returned is the sum of the costs for all possible sets, weighted by the probability of their feasibility being revealed if an agreement is pursued.*
At the function call: $costRound = u_{d(1,1)} \times 0.5 = 1 \times 0.5 = 0.5$.
At the next recursion: $costRound = (u_{d(1,1)} + u_{d(1,2)}) \times 0.25 = (1+2) \times 0.25 = 0.75$.
At the last recursion: $costRound = (u_{d(1,1)} + u_{d(1,2)} + u_{d(1,3)}) \times 0.25 = (1+2+4) \times 0.25 = 1.75$. *The algorithm returns the sum of these three values: estimatedCost=* $0.5 + 0.75 + 1.75 = 3$.
The expected utility of pursuing a solution being positive ($reward-$*estimatedCost=* $4 - 3 = 1$*), the first value is proposed.*

Next is an illustrative example of other *ABTU* operations with the scenario $\{1,3,2\}$.

Example 5. *With the original ABT, A_2 proposes $x_2 = 1$ in Message M_2 and $x_2 = 3$ in Message M_6. In this case, the privacy loss for A_2 is $u_{d(2,1)} + u_{d(2,3)} = 1 + 4 = 5$. However, with ABTU, we do not only use the actual utility of the next assignment to be revealed, but we estimate privacy loss using Algorithms 1 and 2. After A_2 has already sent $x_2 = 1$ with M_2, it considers sending $x_2 = 3$ with M_6. This decision making process is depicted in Figure 2. If the next value, 2 pm, is accepted, A_2 will reach the final state while having revealed $x_2 = 1$ and $x_2 = 3$, for a total privacy cost of $u_{d(2,1)} + u_{d(2,3)} = 1 + 4 = 5$. If it is not, the unavailability of the last value $x_2 = 2$ will have to be revealed to continue the search process, leading to the revelation of all its assignments for a total cost of 7. Since both these scenarios have a probability of 50% to occur, $estimatedCost$ equals $(5+7)/2 = 6$. The utility (*$reward-$*estimatedCost) being equal to $5 - 6 = -1$, A_2 has no interest in revealing $x_2 = 3$ and interrupts the solving. Its final privacy loss is only $u_{d(2,1)} = 2$. The utility of the final state reached by A_2 being -2 with ABTU, and -4 with ABT, ABTU preserves more privacy than ABT in this problem.*

2) Privacy of Assignments: We would like illustrate how our approach deals with privacy of assignment in *UDCOP* by extending the previous example and solving it with *ADOPTU*.

Example 6. *DCOP in Example 2 is extended to a UDCOP by specifying the additional parameters $\mathscr{U}_a$ and $\mathscr{R}$. $\mathscr{U}_a$ represents the cost for each agent to reveal each assigned value. $\mathscr{R}$ is a default reward (for example,* 1000*) that each agent gets for finding a solution to the problem, motivating them to initiate the solving. The description is as follows:*

$$\mathscr{U}_a = \begin{bmatrix} 80 & 20 & 40 \\ 130 & 30 & 10 \\ 80 & 30 & 10 \end{bmatrix} \begin{matrix} A_1 \\ A_2 \\ A_3 \end{matrix} \qquad \mathscr{R} = \{1000, 1000, 1000\}$$

Now we discuss how the standard *DCOP* algorithms are adjusted to *UDCOPs*. After each state transition, each agent computes the estimated utility of the state reached by each possible action, and selects randomly one of the actions leading to the state with the maximum expected utility. In our algorithms, agents estimate expected utilities using the risk of one of their assignments to not be a part of the final solution.

To estimate the cost for a *DCOP* solution, `estimateCostDCOP` is introduced (Algorithm 3). Its inputs are the utilities considered ($\mathscr{U}_i$), the domain of possible values ($\mathscr{D}_{self}$) and the already revealed informations ($revealedInfos$). For each solution, the algorithm evaluates $estimatedCost$ including both constraint costs (*i.e.*, $SolutionCost$ and $PrivacyCost$). Recall that the utility of the reached state is calculated by using Equation 1, and equals the fixed reward (by hypothesis) reduced by $estimatedCost$. If the initial *DCOP* is a maximization problem, it is first recast as a minimization one, so that constraint values correctly belong to costs.

Algorithm 3: `estimateCostDCOP`

Input: $utilities$, $\mathscr{D}_{self}$, $revealedInfos$
Output: $estimatedCost$

```
1 SolutionCost ← 0;
2 PrivacyCost ← 0;
3 foreach value d ∈ D_self do
4     foreach constraint c ∈ C_self do
5         if (c contains predicate p such as p = (x_i = d)) ∧ (d ∈ revealedInfos) then
6             SolutionCost ← SolutionCost + v_{self,p} ;
7             PrivacyCost ← PrivacyCost + utilities_{(self,d)};
8 estimatedCost ← SolutionCost + PrivacyCost;
9 return estimatedCost;
```

Before proposing a new value, agents estimate the utility that will be reached in the next state. This value is the summation of the costs of revealed `agentViews` (weighted by their probability to be the final solution) in the said state, and of the corresponding privacy costs. If this utility is lower than the estimation of the current state, the agent proposes the next value, otherwise it keeps its current value.

a) Asynchronous Distributed Optimization with Utility (ADOPTU): is a method obtained from *ADOPT* by adding Lines 7 to 10 in Algorithm 4 (procedure `checkAgentView_ADOPTU`). At Line 7, the possible next value is set to the value that has the minimal cost. The cost reached after the next value (Line 8) and the cost of the current state (Line 9) are estimated. At Line 10, if the next cost is lower than the current cost, the maximal improvement and the new value are updated.

Example 7. *Continuing with Example 2, at the beginning of the computation with the ADOPTU solver, the participants select a random value. The resulting set of assignments is $x_1 = 1$, $x_2 = 3$ and $x_3 = 2$ for A_1, A_2, and A_3 respectively. The participants then inform their values to their linked descendants (A_2 and A_3 for A_1, and A_3 for A_2).*

A_2 will have `agentView` $= \{(x_1 = 1)\}$ *and chooses the value* 1*, since this value minimizes its cost. A_3 sends a VIEW message with cost of* 190 *to A_2. This cost is reported because given the value of x_1, it is a lower bound on global solution cost:* 190 *is a lower bound on local cost at x_2, and* 0 *is a lower bound on local costs at other variables. Now concurrently with the previous execution of A_2, A_3 receives the assignment $x_2 = 3$. With ADOPTU, A_3 realizes that changing its value would increase its cost and decides to keep its value unchanged. However, with ADOPT, A_3 decides to change its value to* 3*, as it reduces the cost. Then, with ADOPT and ADOPTU, A_3 receives the assignments $x_1 = 1$ and $x_2 = 1$, and decides to change its value to $x_3 = 1$. At the final step, the previous* `agentView` *is the optimal solution. With ADOPTU, the reached costs are: $70 + 80 = 150, 120 + 30 + 100 = 250, 80 + 80 = 160$ for A_1, A_2, and A_3 respectively. With standard ADOPT, the final utilities are: $70 + 80 = 150, 120 + 30 + 100 = 250, 80 + 80 + 30 = 190$ for A_1, A_2, and A_3 respectively. Therefore, by using ADOPTU instead of ADOPT, A_3 avoids a futile revelation of $x_3 = 2$ and reduces its utility by* 30*.*

3) Privacy of Constraints: Similarly, we detail the model for our example and we present the extensions of two solvers for *UDCOPs* (namely *DBOU* and *DSAU*).

Algorithm 4: `checkAgentView_ADOPTU`

Input: $utilities$, $\mathscr{D}_{self}$, $revealedAssignments$
1 **for each** ($d \in \mathscr{D}_{self}$) update $l[d]$ and recompute $h[d]$;
2 **for each** A_j with higher priority than A_{self} **do**;
3 **if** *(h has non-null cost CA for all values of $\mathscr{D}_{self}$)* **then**
4 $vn \leftarrow$ min_resolution(j);
5 **if** $vn \neq lastSent[j]$ **then**
6 send **nogood**(vn,$self$) to A_j;
7 $newValue \leftarrow \text{argmin}_d(cost[h[d]])$;
8 $nextCost \leftarrow$ `estimateCostDCOP` ($utilities$, $\mathscr{D}_{self}$, $revealedAssignments \cup newValue$);
9 $currentCost \leftarrow$ `estimateCostDCOP` ($utilities$, $\mathscr{D}_{self}$, $revealedAssignments$);
10 **if** *($nextCost < currentCost$)* **then**
11 assign $newValue$ to x_{self} ;
12 **if** *(x_{self} was modified)* **then**
13 send **ok?** to each neighbour to inform them about the value change;

Example 8. *DCOP described in Example 2 is extended to a UDCOP by specifying the parameters $\mathscr{C}$, $\mathscr{U}_c$ and $\mathscr{R}$ as follows. $\mathscr{C}$ is the updated set of constraints. $\mathscr{U}_c$ represents the cost for each agent to reveal each assignment cost. $\mathscr{R}$ is a default reward that each agent gets for finding a solution to the problem, motivating them to initiate the solving. $\mathscr{C}$'s subsets known to each agent are:*

$\mathscr{C}_1 = \{c_{1,1} = [(x_1 = 1), 70], c_{1,2} = [(x_1 = 2), 230],$
$c_{1,3} = [(x_1 = 3), 270], c_{1,4} = [\neg(x_1 = x_2 = x_3), \infty]\}$
$\mathscr{C}_2 = \{c_{2,1} = [(x_2 = 1), 120], c_{2,2} = [(x_2 = 2), 400],$
$c_{2,3} = [(x_2 = 3), 190], c_{2,4} = [\neg(x_1 = x_2 = x_3), \infty]\}$
$\mathscr{C}_3 = \{c_{3,1} = [(x_3 = 1), 40], c_{3,2} = [(x_3 = 2), 280],$
$c_{3,3} = [(x_3 = 3), 230], c_{3,4} = [\neg(x_1 = x_2 = x_3), \infty]\}$

$$\mathscr{U}_c = \begin{bmatrix} 80 & 20 & 40 \\ 100 & 30 & 10 \\ 80 & 30 & 10 \end{bmatrix} \begin{matrix} A_1 \\ A_2 \\ A_3 \end{matrix} \qquad \mathscr{R} = \{1000, 1000, 1000\}$$

To adapt existing algorithms for privacy of constraints, revealed domains and possible revealed domains are changed to the revealed constraints and possible revealed constraints, respectively.

a) Distributed Breakout with Utility (DBOU): is a solver obtained from *DBO* by adding Lines 5 to 11 in Algorithm 5 (procedure `sendImproveDBOU`). At each iteration, *DBOU* does not only uses solution cost to guide search and computes the value giving maximal improvement as in standard *DBO* (Lines 2 to 4), but also considers constraint privacy costs (Lines 5 to 7). The new chosen value is the one minimizing total cost (Lines 8 to 11). As privacy loss is cumulative, agents update the set of revealed constraints to also consider previously revealed constraints during their estimation of the reached cost for the different considered values.

b) Distributed Stochastic Algorithm with Utility (DSAU): is obtained from standard *DSA* by adding Lines 10 to 13 in Algorithm 6 (procedure `DSAU`). Each agent computes a solution and sends it to its neighbours (Lines 1 to 5). At each iteration, after collecting new values from neighbours (Line 6), each agent compares the new `agentView` with the previous one (Line 7). If a difference is detected, the agent A_{self} computes a new solution considering both solution (Lines 8 and 9) and privacy costs (Lines 10 to 13) similarly with *ADOPTU* and *DBOU*.

Example 9. *Continuing with Example 2, at the beginning of the computation with the DSAU solver, the participants select a random value. The resulting* `agentView` *of each agent is :*
$\{(x_1 = 1), (x_2 = 1), (x_3 = 3)\}$. *The utilities of the reached state are:*
$v_{1,1} + u_{c(1,1)} = 70 + 80 = 150$,
$v_{2,1} + u_{c(2,1)} = 120 + 100 = 220$, *and*
$v_{3,3} + u_{c(3,3)} = 230 + 10 = 240$ *for Students A_1, A_2, and A_3 respectively.*
The participants then inform each other of their value. They consider changing them to a new randomly selected one. The considered `agentView` *is* $\{(x_1 = 2), (x_2 = 3), (x_3 = 1)\}$. *If the participants change their value, the utilities of the reached states would be:*
$(v_{1,1} + v_{1,2})/2 + u_{c(1,1)} + u_{c(1,2)} = 250$,
$(v_{2,1} + v_{2,3})/2 + u_{c(2,1)} + u_{c(2,3)} = 265$,

Algorithm 5: `sendImproveDBOU`

Input: $utilities$, $\mathscr{D}_{self}$, $revealedConstraints$

$eval \leftarrow$ evaluation value of x_{self};
$myImprove \leftarrow 0$;
$newValue \leftarrow x_{self}$;
$nextValue \leftarrow$ value giving maximal improvement;
$nextRevealedConstraints \leftarrow revealedConstraints \cup$ constraints with $nextValue$;
$nextCost \leftarrow$ `estimateCostDCOP` ($utilities$, $\mathscr{D}_{self}$, $nextRevealedConstraints$);
$currentCost \leftarrow$ `estimateCostDCOP` ($utilities$, $\mathscr{D}_{self}$, $revealedConstraints$);
if *($nextCost < currentCost$)* **then**
    $myImprove \leftarrow$ possible max improvement;
    $newValue \leftarrow$ value giving maximal improvement;
    $x_{self} \leftarrow newValue$;
if *($eval = 0$)* **then**
    $consistent \leftarrow true$
else
    $consistent \leftarrow false$;
$terminationCounter \leftarrow 0$;
if *($myImprove > 0$)* **then**
    $canMove \leftarrow true$;
    $quasiLocalMin \leftarrow false$;
else
    $canMove \leftarrow false$;
    $quasiLocalMin \leftarrow true$
send (**improve**, x_{self}, $myImprove$, $eval$, $terminationCounter$) to neighbours;

Algorithm 6: DSAU

Input: $utilities$, $\mathscr{D}_{self}$, $revealedConstraints$

$nextValue \leftarrow$ randomly chosen a value;
while *(no termination condition is met)* **do**
    if *($nextValue \neq x_{self}$)* **then**
        $x_{self} \leftarrow nextValue$;
        send x_{self} to neighbours;
    $nextView \leftarrow$ collect x_i for each neighbour A_i;
    if *($nextView \neq$ `agentView`)* **then**
        `agentView`$\leftarrow nextView$;
        $tempValue \leftarrow$ randomly chosen value;
        $revealedConstraints \leftarrow revealedConstraints \cup \mathscr{C}_{self}$ with $x_{self} = tempValue$;
        $nextCost \leftarrow$ `estimateCostDCOP` ($utilities$, $\mathscr{D}_{self}$, $nextRevealedConstraints$);
        $currentCost \leftarrow$ `estimateCostDCOP` ($utilities$, $\mathscr{D}_{self}$, $revealedConstraints$);
        if *($nextCost < currentCost$)* **then**
            $x_{self} \leftarrow tempValue$;

$(v_{3,3} + v_{3,1})/2 + u_{c(3,3)} + u_{c(3,1)} = 225$*, for Students* A_1*,* A_2*, and* A_3 *respectively.*
A_1 *and* A_3 *do not propose the new value as it would increase their utility. However,* A_3 *chooses to change its value from* 2 *to* 1 *which lowers its utility from* 240 *to* 225*. In the next step,* `agentView` *is* $\{(x_1 = 1), (x_2 = 1), (x_3 = 1)\}$*. Participants then do not change their value any-more, as all other options would not decrease the utility. At the final step, the previous* `agentView` *is therefore the optimal solution. With DSAU, the reached utilities are :*
$70 + 80 = 150, 120 + 100 = 220, 40 + 10 + 80 = 130$ *for Students* A_1*,* A_2*, and* A_3 *respectively. With standard DSA, the final utilities are:*
$(v_{1,1} + u_{c(1,1)} + u_{c(1,2)} + u_{c(1,3)}) = 230$*,*
$(v_{2,1} + u_{c(2,1)} + u_{c(2,2)} + u_{c(2,3)}) = 260$*,*
$(v_{3,1} + u_{c(3,2)} + u_{c(3,1)} + u_{c(3,3)}) = 160$*, for Students* A_1*,* A_2*, and* A_3 *respectively.*

Therefore, using DSAU instead of DSA reduces the utility by $80, 40, 30$.

In this work, studied problems include only one type of privac at a time, to illustrate proposed models and algorithms with simple examples. However, problems integrating several types of privacy can also be modelled with *UDisCP*. Such problems, where agents would have optimize multiple objectives, will be investigated in future works.

IV. Theoretical Discussion

This section deals with three theoretical studies, *i.e.*, comparisons with *DCOPs* (IV-A), *MO-DCOPs* (IV-B) and *POMDP* (IV-C).

A. Comparison with DCOPs

The introduced *UDCOP* framework can assume that inter-agent constraints are public (without significant loss of generality). This is due to the fact that any problem with private inter-agent constraints, is equivalent with its dual representation where each constraint becomes a variable [3].

Theorem 1. *UDCOP planning and execution is at least as general as DCOPs solving.*

Proof. A *DCOPs* can be modelled as a *UDCOPs* with all privacy costs equal 0. The obtained *UDCOPs* would always reach an agreement, if possible. Therefore the goal of a *UDCOPs* would also match with the goal of the modelled *DCOPs*. This implies a tougher class of complexity for *UDCOPs*. □

The space complexity required by *ABTU* and *SyncBTU* in each agent is identical with the one of *ABT* and *SyncBT*, since the only additional structures are the privacy costs associated with its values, constituting a constant factor increases for domain storage Similarly, additional structures with constant values are added from *DCOP* to *UDCOP*.

B. Comparison with MO-DCOPs

a) Multi-Objective Distributed Constraint Optimization Problem: A multi-objective optimization problem (*MOOP*) [37], [11] is defined as the problem of simultaneously maximizing k objective functions that have no common measure, defined over a set of variables, each one taking its value in a given domain. Thus, a solution to *MOOP* is a set of assignments maximizing the combination of the objective functions. Here, each objective function can be defined over a subset of variables of the problem. However, to simplify our discourse, we assume that each function is defined over the same set of variables. A Multi-Objective *DCOP* [8], called *MO-DCOP*, is an extension of the standard mono-objective *DCOPs*.

Note that a *MO-DCOP* is a *DCOP* where the weight of each constraint tuple is a vector of values $[w_i]$, each value w_i representing a different metric. Two weights $[w_i^1]$ and $[w_i^2]$ for the same partial solution, inferred from disjoint sets of weighted constraints, are combined into a new vector $[w_i^3]$ where each value is obtained by summing the values in the corresponding position two input vectors, namely $w_i^3 = w_i^1 + w_i^2$. The quality of a solution of *MO-DCOP* is a vector integrating the cost of all weighted constraints. The vectors can be compared using various criteria, such as leximin, maximin, social welfare or Theil index [45], [41].

To clarify why Multi-Objective *DCOPs* (*MO-DCOPs*) cannot integrate our concept of privacy as one of the criteria they aggregate, we give an example of what would be achieved with *MO-DCOPs*, as contrasted with the results using *DCOPs*. We show a comparative trace based on one of the potential techniques in *MO-DCOPs*, providing a hint on why *MO-DCOPs* cannot aggregate privacy lost during execution in the same way as *UDCOP*. In this example, the privacy value of each assignment and its constraint cost are two elements of an ordered pair defining the weight of *MO-DCOP*. For illustration, in this example pairs of weights are compared lexicographically with the privacy having priority.

Example 10. *We assume to model Example 2 with a MO-DCOP. As also illustrated in the trace (Table I) with lexicographical comparison (UDCOP DSAU vs. MO-DCOP DSA), privacy first. Candidate values are marked with * if they are better than old values, and will be adopted. For these two approaches, we evaluate both* $SolutionCost$ *and* $PrivacyCost$.

At the beginning of DSA process, the participants select a random value. Resulting `agentView` *is* $\{(x_1 = 1), (x_2 = 1), (x_3 = 3)\}$. *For instance,* A_1*evaluates this state with* $PrivacyCost$ *of* 80 *and* $SolutionCost$ *of* 70. *Recall that privacy criterion is prevailing; hence a notation* $[80, 70]$. *DSAU evaluates this same state as a summation of these costs, with a value of* 150. *The participants then inform each others of their value. They consider a change of their value to a new randomly selected one.* `agentView` *is* $\{(x_1 = 2), (x_2 = 3), (x_3 = 1)\}$ *for these two solvers.*

With UDCOP DSAU (state$_1$*),* A_1*and* A_2 *do not propose the new value as it would increase their cost, and* A_3 *chooses to change its variable's value from* 3 *to* 1. *However, with MO-DCOP DSA,* A_2 *changes its value to* 3 *(as it believed* $PrivacyCost$ *will drop from* 100 *to* 10*), which is not the case with DSAU, which implies privacy loss. The* `agentView` *is now* $\{(x_1 = 2), (x_2 = 3), (x_3 = 3)\}$ *(the real privacy loss is the summation of all revealed costs, i.e.,* $100 + 10 = 110$*, and not only* 10*).*

In short, with the MO-DCOP model, A_2 *reveals more values and loses more privacy (with a difference of* $110 - 100 = 10$ *for privacy costs) than with UDCOPs.*

TABLE I
COMPARATIVE TRACE OF TWO ROUNDS (*UDCOP DSAU* VS. *MO-DCOP DSA*)

Framework	*UDCOP DSAU*			*MO-DCOP DSA*		
Agent	A_1	A_2	A_3	A_1	A_2	A_3
	current state					
state_0	1	1	3	1	1	3
$SolutionCost$	70	120	230	70	120	230
$PrivacyCost$	80	100	10	80	100	10
evaluation	150	220	240	$[80, 70]$	$[100, 120]$	$[10, 230]$
	believed next state					
considered state	2	3	1	2	3	1
$SolutionCost$	150	155	135	230	190	40
$PrivacyCost$	100	110	90	20	10	80
evaluation	250	265	225*	$[20, 230]$*	$[10, 190]$*	$[80, 40]$
	achieved next state					
state_1	1	1	1	2	3	3
$SolutionCost$	70	120	40	230	190	230
$PrivacyCost$	80	100	90	100	110	10
evaluation	150	220	130	$[100, 230]$	$[110, 190]$	$[10, 230]$

The assumption that each agent owns a single variable is also not restrictive. Multiple variables in an agent can be aggregated into a single variable by Cartesian product. Nevertheless some algorithms can exploit these underlying structures for efficiency, and this has been the subject of extensive research [17], [40].

We now discuss how *UDisCP* can be interpreted as a planning problem.

C. Comparison with POMDPs

The problem that each agent in *UDisCP* has to solve have similarities to a Partially Observable Markov Decision Problem (*POMDP*). Given ways to approximate observation and transition conditional probabilities, these problems could be reduced to *POMDPs* [43], [30], [12].

A *POMDP* agent regularly reasons in terms of belief (probability distribution over the states), and tries to build a policy, namely a recommendation of each action to be executed as function of current belief. For *UDisCP*, the corresponding *POMDP* is defined by the tuple $\langle S, A, T, R, \Omega, O, \gamma \rangle$ with components [50]:

- $\mathcal{S}$: Set of states of the agent, defined by possible contents of its `agentView`, of the nogoods stored by the agent, the knowledge the agent gathers about the secret elements of the *UDisCP* unknown to it, and the information already revealed.
- $\mathcal{A}$: Set of actions available to an agent, consisting in local reasoning and communication actions that are a function of the selected protocols (*i.e.*, communication language). For example, in *ABT* these communications actions can have as payloads assignment announcements (`ok?` messages) and nogoods in (`nogood` messages).
- T: Set of transition probabilities between states given actions for *UDisCP*. It is estimated in our approach by training $agreementProb$.
- $\mathcal{R}$: Set of rewards of *POMDP* is the same as $\mathscr{R}$ for the corresponding *UDisCP*.
- Ω: Set of possible observations is given by $\mathscr{U}$, the possible incoming payloads of the communication actions available in *UDisCP*, as well as possible results of local reasoning steps.
- O: Set of conditional observation probabilities. In reported experiments, it is assumed that the message payloads truthfully reveal the corresponding elements of the states of the other agents, while the probability of the remaining elements have to be inferred by the agent.
- γ: Discount factor set to 1, since we have not taken into consideration the impact of time on utilities in this work.

In the next section, we will present experiments that use our *UDisCP* models and algorithms to preserve privacy.

V. EXPERIMENTAL RESULTS

We describe the experimental protocol for our study (V-A), and then we give the obtained results for DMS context (V-B).

A. Experimental protocol and DMS context

We evaluate our framework and algorithms on randomly generated instances of *distributed meeting scheduling problems (DMS)* [37], [5], [19]. Existing studies has already addressed the question of privacy in DMS problems by considering the information on whether an agent can attend a meeting to be private [59], [7].

The algorithm we use to generate the DMS problems is defined according to the following procedure: (i) Variables are created and associated with the agent controlling them. (ii) Domains (possible values) are initialized for all variables. (iii) The global constraint *"all equals"* is added. (iv) The unary constraints (individual unavailabilities) are added. (v) Privacy costs uniformly distributed between 0 and 9 are generated for each possible variable value, variable assignment and constraint.

The experiments are carried out on a computer under Windows 7, using a 1 core 2.16 GHz CPU and 4 GiB of RAM. Implementation is done in Java (jre 1.8) using the JADE platform (version 4.3.3) to build the multi-agent system [4].

The problems are parametrized as follows: 10, 20, or 40 agents, 10 possible values, and the cost of a revelation is a random number between 0 and 9. These parameters are used to guarantee that the problems are not over-constrained or under-constrained, as the probability to find a solution increases with the increase of domain size, and with the decrease of the number of agents and unary constraint tightness. Density is defined as the proportion of binary constraints in $\mathscr{C}$ among $\mathscr{X}_i \times \mathscr{X}_j$ $(i \neq j)$. Because of the constraint requiring all couples of variables to be equals, density in DMS is always 1. Tightness is defined as proportion of forbidden tuples among all constraints (binary but also unary ones) of $\mathscr{C}$.

Proof. For DMS, we assume t for constraint tightness, m for the number of agents $\mathscr{A}$ and d for the size of each domain set $\mathscr{D}_i$. Indeed, the probability that a given value is authorized is:

$$1 - t$$

The probability that a given value is authorized by all agents is:

$$(1-t)^m$$

The probability that a given value is forbidden by at least one agent is:

$$1 - (1-t)^m$$

The problem has a probability s to have at least one solution if and only if:

$$1 - s = (1 - (1-t)^m)^d$$

$$1 - \sqrt[d]{(1-s)} = (1-t)^m$$

$$1 - t = \sqrt[m]{(1 - \sqrt[d]{(1-s)})}$$

□

Figures 3 corroborates this for a DMS modelled by *DisCSP* with a probability of 50% to have at least one solution, as depicted according to the previous formula and to experimental results. These figures are described by the number of agents (coordinate x), domain size (coordinate y) and tightness (coordinate z).

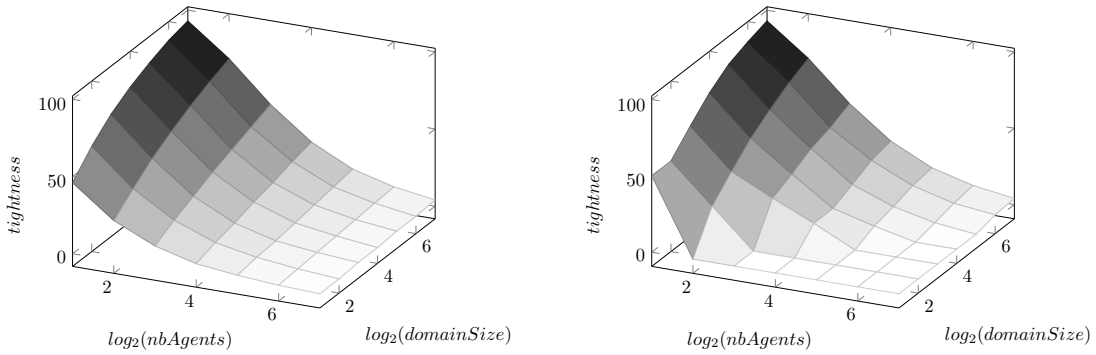

Fig. 3. Parametrization for a probability of 50% to have a solution – theoretical (left) and experimental (right) values

For example, to have a probability of 50% to have a solution, we see that a DMS with 8 agents ($x = 3$), and a domain size of 8 ($y = 3$), needs to be parametrized with a tightness of 27% ($z = 27$). For a problem with 16 agents (*i.e.*, $x = 4$) and a same domain size ($y = 3$) and a probability to have a solution, the tightness becomes 14% ($z = 14$). Note that in the experimental plot, values may slightly differ due to the discrete nature of domains size. Later, this formula will be used to generate problems with relevant parameters values and that are not under- or over-constrained. We will verify that probability to find a solution to DMS is negatively correlated with the number of agents, constraints tightness, and positively correlated with domains size. Each set of reported experiments is an average estimation of 50 instances for the different algorithms (*SyncBT – SyncBTU*, *ABT – ABTU*, *ADOPT – ADOPTU*, *DBO – DBOU* and *DSA – DSAU*).

B. Experiments on UDisCP

Tables II and III show the average privacy loss per agent during the execution of the solving algorithms for *DisCSPs* (*SyncBT*, *ABT*), *DCOPs* (*DBO*, *DSA*, and *ADOPT*), and of their respective extensions, with several values for the number of agents, the domain size and the constraints tightness. Bold data points show instances with high privacy loss (from 10.0 to 20.0). Privacy loss of 20.0 is the maximal value before interrupting solving. Empty data points show instances with low privacy loss (below 0.1). Moreover, we refer to a data point in these tables as $solver(nbAgents, domainSize, tightness)$. For example, *SyncBT*(10,20,30) refers to the average privacy loss for instances with 10 agents, a domain size of 20, and a constraint tightness of 30% solved with *SyncBT*, namely 6.7.

TABLE II
PRIVACY LOSS IN *DisCSPs* AND *UDisCSPs*

nbAgents	10																			
domainSize	10					20					30					40				
tightness(%)	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50
SyncBT	0.7	3.7	3.8	3.3	3.5	0.7	4.9	6.7	6.0	7.0	0.7	5.2	8.8	9.3	**10.5**	0.7	5.4	**10.3**	**12.1**	**14.0**
SyncBTU	0.7	3.4	2.8	2.0	0.6	0.7	4.5	4.9	3.7	1.2	0.7	4.8	6.5	5.7	1.8	0.7	4.9	7.9	7.3	2.4
ABT	3.1	7.2	**12.4**	**13.9**	1.6	3.1	9.5	**20.0**	**20.0**	3.2	3.1	**10.3**	**20.0**	**20.0**	4.8	3.1	**10.4**	**20.0**	**20.0**	6.4
ABTU	2.5	5.2	6.1	6.9	1.5	2.5	6.9	**10.7**	**11.3**	3.0	2.5	7.4	**14.2**	**19.5**	4.5	2.5	7.5	**16.6**	**20.0**	6.0
nbAgents	20																			
domainSize	10					20					30					40				
tightness(%)	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50
SyncBT	0.5	0.4	0.1			0.7	1.1	0.1			0.7	4.7	0.1			0.7	1.9	0.1		
SyncBTU	0.5	0.3	0.1			0.7	1.0	0.1			0.7	4.3	0.1			0.7	1.8	0.1		
ABT	2.3	0.7	0.1			2.9	2.1	0.1			3.0	9.3	0.1			3.1	3.8	0.1		
ABTU	1.8	0.5	0.1			2.3	1.5	0.1			2.4	6.7	0.1			2.5	2.8	0.1		
nbAgents	40																			
domainSize	10					20					30					40				
tightness(%)	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50
SyncBT	0.1					0.2					0.3					0.3	0.1			
SyncBTU	0.1					0.2					0.3					0.3	0.1			
ABT	0.4					0.8					0.9					1.1	0.1			
ABTU	0.4					0.7					0.9					1.1	0.1			

TABLE III
PRIVACY LOSS IN *DCOPs* AND *UDCOPs*

nbAgents	10																			
domainSize	10					20					30					40				
tightness(%)	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50
DBO	0.3	0.6	1.1	1.1	1.6	0.4	0.8	1.4	1.4	2.1	0.6	1.1	2.0	2.0	2.9	0.8	1.5	2.8	2.8	4.1
DBOU	0.2	0.4	0.4	0.8	1.0	0.3	0.5	0.5	1.0	1.3	0.4	0.7	0.7	1.4	1.8	0.5	1.0	1.0	2.0	2.6
DSA	1.4	3.6	5.0	6.4	8.4	1.8	4.7	6.5	8.3	**10.9**	2.9	6.6	9.1	**11.7**	**15.3**	3.6	9.2	**12.8**	**16.3**	**20.0**
DSAU	1.4	2.0	2.2	2.6	3.0	1.8	2.6	2.9	3.4	3.9	2.6	3.6	4.0	4.7	5.5	3.6	5.1	5.6	6.6	7.6
ADOPT	3.6	4.2	5.6	6.8	8.8	4.4	5.5	7.3	8.8	**11.4**	6.2	7.6	**10.2**	**12.4**	**16.0**	8.7	1.7	**14.2**	**17.3**	**20.0**
ADOPTU	2.4	3.0	3.4	4.0	4.2	3.1	3.9	4.4	5.2	5.5	4.4	5.5	6.2	7.3	7.6	6.2	7.6	8.7	**10.2**	**10.7**
nbAgents	20																			
domainSize	10					20					30					40				
tightness(%)	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50
DBO	0.2	0.1	0.1			0.4	0.7	0.1			0.5	1.0	0.1			0.7	1.4	0.1		
DBOU	0.2	0.1	0.1			0.2	0.1	0.1			0.3	0.7	0.1			0.5	0.4	0.1		
DSA	1.0	0.4	0.1			1.6	1.0	0.1			2.5	5.9	0.1			3.2	3.3	0.1		
DSAU	1.0	0.2	0.1			1.6	0.6	0.1			2.5	3.3	0.1			3.4	1.8	0.1		
ADOPT	2.5	0.5	0.1			4.0	1.2	0.1			6.0	6.9	0.1			8.3	0.6	0.1		
ADOPTU	1.8	0.3	0.1			2.8	0.2	0.1			4.2	4.9	0.1			5.9	2.8	0.1		
nbAgents	40																			
domainSize	10					20					30					40				
tightness(%)	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50	10	20	30	40	50
DBO	0.1					0.1					0.6					0.8	0.1			
DBOU	0.1					0.1					0.1					0.2	0.1			
DSA	0.3					0.5					1.0					1.3	0.1			
DSAU	0.3					0.4					1.0					1.3	0.1			
ADOPT	0.8					1.2					2.5					3.0	0.1			
ADOPTU	0.5					0.8					1.8					2.2	0.1			

a) : Generally, we see that for both Tables II and III, increasing the number of agents implies a reduction of the number of solutions. Thus, problems being over-constrained, agents interrupt the solving faster, which can explain the reduction of average privacy loss for problems with many agents. Indeed, we see that most instances with 10 agents have a significant average privacy loss, while instances with 40 agents have low privacy loss. For example, in Table II, for $solver$(10, 10, 50), all solvers imply a privacy loss (3.5, 0.6, 1.6, 1.5) for *SyncBT* – *SyncBTU* and *ABT* – *ABTU*, respectively. However, for $solver$(40, 10, 50) all the previous solvers have empty data points. Similarly for Table III, results obtained for 40 agents give also a lower privacy losses for the same reasons.

b) : Moreover, we also see that for these two tables, high privacy loss is correlated with a high tightness and a high domain size. Table II shows that the number of high privacy loss depend on the tightness (more than 30%) and a domain size up to 30. For example, for $ABTU$(10,10,50), privacy loss is 1.5, while it is 6.0 for *ABTU*(10,40,50). Table III shows that

high privacy loss (bold data points) never occurs for instances with a domain size of 10, while it does for most problems with a domain size of 40. For example, for *ADOPT*(10,10,50), privacy loss is 8.8, while *ADOPT*(10,40,50) gives 20.0. For these two tables, the main reason is that a higher domain size allows more solution proposals and avoids premature interruptions.

c) : Finally, recall that the difference of values between different families of solvers is explained by the different types of privacy considered (*i.e.*, domain, assignment and constraint), as the number of revelations of assignments differs from the number of revelations of constraints. Table II measures domain privacy loss for agents during the execution of *SyncBT*, *ABT*, and their extensions. *SyncBT* and *SyncBTU* are better than *ABT* and *ABTU* at preserving privacy, likely due to the increase of exchanged messages with asynchronous solvers, as agents can run concurrently. Table III measures assignment privacy with *ADOPT* and *ADOPTU*. For *ADOPT*(10,40,t) with $t \in \{10, 20, 30, 40, 50\}$, the assignment privacy loss is $\{6.2, 7.6, 10.2, 12.4, 16\}$. However with the same parameters for *ADOPTU*, it is $\{4.4, 5.5, 6.2, 7.3, 7.6\}$. The loss is reduced with our extension as search is driven by objective to minimize the assignment privacy loss, and interrupt solving if necessary. Table III measures constraint privacy loss with *DBO – DBOU* and *DSA – DSAU*. We see that *DBO* and *DBOU* are better than *DSA* and *DSAU*, likely due to the fact that with *DBO*, only one agent changes its value at each iteration, while with *DSA* all agents do, which implies an increase of the number of exchanged messages. For example, in Table III for $solver(10,\ 20,\ 50)$ privacy loss drops from 2.1 to 1.3 and from 10.9 to 3.9 for *DBO – DBOU* and *DSA – DSAU* respectively. For all these algorithms dealing with different types of privacy, an increase of privacy loss with initial solvers implies a better preservation with extended ones.

VI. Conclusion

While many approaches have been proposed recently for dealing with privacy in distributed constraint satisfaction and optimization problems, none of them is exempt from limitations. These approaches may require particular properties from the initial problem, or may consider certain aspects of privacy only. In this work, we propose the Utilitarian Distributed Constrained Problem (*UDisCP*) framework. The framework models the privacy loss for the revelation of an agent's constraints as a utility function, letting agents integrate privacy requirements directly in their search process. Solving the problem then consists in finding the best compromise between solution quality and privacy loss, instead of focusing only on solution quality. We propose extensions to existing algorithms for *DisCSPs* (*SyncBTU*, *ABTU*) and *DCOPs* (*DBOU*, *DSAU*, and *ADOPTU*) that let agents use information about privacy to modify their behaviour and guide their search process, by proposing values that reduce the amount of privacy loss, and compare them on different types of distributed meeting scheduling problems. The comparison shows that explicit modelling and reasoning with the utility of privacy allows for significant savings in privacy with minimal impact on the quality of the achieved solutions.

Our approach has several possible extensions. For future works, we first plan to extend our models and algorithms to more general problems modelled with constraints, namely problems with n-ary constraints (not only binary ones) as well as multi-variable problems, where each agent may control several variables. In this case, the solution of an agent is a set of assignments, rather than a single assignment. Further, we want to investigate the notion of ethics [9], implying that agents may have remorse for lying when modifying their behaviour for privacy. These notion of ethics leads to the building of communities of interests between agents [65], [66], where privacy costs will differ according to the recipient of a message, instead of being random. A *DisCSP* modelling was used for road traffic in [13], [14], and we would like to apply our model of privacy for road traffic simulation where the agents/drivers may choose whether to reveal some private data concerning for example their driving, or their habits. Indeed, we think that our privacy model may also be adapted to other applications such as multi-robot exploration [44] or smart energy [49].

Acknowledgements

Authors gratefully acknowledge the reviewers for their constructive remarks, allowing to improve our article. This work was realized while the first author was a visiting researcher at Florida Institute of Technology (Sept. 2015 – May 2016).

References

[1] Alejandro Arbelaez, Youssef Hamadi, and Michèle Sebag. Continuous search in constraint programming. In *22nd IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2010, Arras, France, 27-29 October 2010 - Volume 1*, pages 53–60. IEEE Computer Society, 2010.

[2] Alejandro Arbelaez, Youssef Hamadi, and Michele Sebag. Continuous search in constraint programming. In Youssef Hamadi, Eric Monfroy, and Frédéric Saubion, editors, *Autonomous Search*, chapter 9, pages 219–243. Springer Berlin Heidelberg, 2012.

[3] Fahiem Bacchus and Peter Van Beek. On the conversion between non-binary and binary constraint satisfaction problems. In Jack Mostow and Chuck Rich, editors, *Proceedings of the Fifteenth National Conference on Artificial Intelligence and Tenth Innovative Applications of Artificial Intelligence Conference, AAAI 98, IAAI 98, July 26-30, 1998, Madison, Wisconsin, USA.*, pages 310–318. AAAI Press / The MIT Press, 1998.

[4] Fabio Bellifemine, Federico Bergenti, Giovanni Caire, and Agostino Poggi. JADE - A java agent development framework. In Rafael H. Bordini, Mehdi Dastani, Jürgen Dix, and Amal El Fallah-Seghrouchni, editors, *Multi-Agent Programming: Languages, Platforms and Applications*, volume 15 of *Multiagent Systems, Artificial Societies, and Simulated Organizations*, pages 125–147. Springer, 2005.

[5] Christian Bessiere, Emmanuel Hebrard, Brahim Hnich, and Toby Walsh. The complexity of reasoning with global constraints. *Constraints*, 12(2):239–259, 2007.

[6] Ismel Brito, Amnon Meisels, Pedro Meseguer, and Roie Zivan. Distributed constraint satisfaction with partially known constraints. *Constraints*, 14(2):199–234, 2009.

[7] Ismel Brito and Pedro Meseguer. Privacy in distributed meeting scheduling. In Teresa Alsinet, Josep Puyol-Gruart, and Carme Torras, editors, *Artificial Intelligence Research and Development, Proceedings of the 11th International Conference of the Catalan Association for Artificial Intelligence, CCIA 2008, October 22-24, 2008, Sant Martí d'Empúries, Spain*, volume 184 of *Frontiers in Artificial Intelligence and Applications*, pages 118–127. IOS Press, 2008.

[8] Maxime Clement, Tenda Okimoto, Tony Ribeiro, and Katsumi Inoue. Model and algorithm for dynamic multi-objective distributed optimization. In Guido Boella, Edith Elkind, Bastin Tony Roy Savarimuthu, Frank Dignum, and Martin K. Purvis, editors, *International Conference on Principles and Practice of Multi-Agent Systems*, volume 8291 of *Lecture Notes in Computer Science*, pages 413–420. Springer, 2013.

[9] Nicolas Cointe, Grégory Bonnet, and Olivier Boissier. Ethical judgment of agents' behaviors in multi-agent systems. In Catholijn M. Jonker, Stacy Marsella, John Thangarajah, and Karl Tuyls, editors, *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, pages 1106–1114. International Foundation for Autonomous Agents and Multiagent Systems, ACM, 2016.

[10] Ludivine Crépin, Yves Demazeau, Olivier Boissier, and François Jacquenet. Privacy preservation in a decentralized calendar system. In Yves Demazeau, Juan Pavón, Juan M. Corchado, and Javier Bajo, editors, *7th International Conference on Practical Applications of Agents and Multi-Agent Systems, PAAMS 2009, Salamanca, Spain, 25-27 March 2009*, volume 55 of *Advances in Intelligent and Soft Computing*, pages 529–537. Springer, 2009.

[11] Francesco M Delle Fave, Ruben Stranders, Alex Rogers, and Nicholas R Jennings. Bounded decentralised coordination over multiple objectives. In Liz Sonenberg, Peter Stone, Kagan Tumer, and Pinar Yolum, editors, *10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011), Taipei, Taiwan, May 2-6, Volume 1-3*, pages 371–378. IFAAMAS, 2011.

[12] Jilles Steeve Dibangoye, Guy Shani, Brahim Chaib-draa, and Abdel-Illah Mouaddib. Topological order planner for pomdps. In Craig Boutilier, editor, *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, pages 1684–1689, 2009.

[13] Arnaud Doniec, René Mandiau, Sylvain Piechowiak, and Stéphane Espié. Anticipation based on constraint processing in a multi-agent context. *Autonomous Agents and Multi-Agent Systems*, 17(2):339–361, 2008.

[14] Arnaud Doniec, René Mandiau, Sylvain Piechowiak, and Stéphane Espié. Controlling non-normative behaviors by anticipation for autonomous agents. *Web Intelligence and Agent Systems*, 6(1):29–42, 2008.

[15] Prashant Doshi, Toshihiro Matsui, Marius Silaghi, Makoto Yokoo, and Markus Zanker. Distributed private constraint optimization. In *Proceedings of the 2008 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Sydney, NSW, Australia, December 9-12, 2008*, pages 277–281. IEEE Computer Society, 2008.

[16] Boi Faltings, Thomas Léauté, and Adrian Petcu. Privacy guarantees through distributed constraint satisfaction. In *Proceedings of the 2008 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Sydney, NSW, Australia, December 9-12, 2008*, volume 2, pages 350–358. IEEE Computer Society, 2008.

[17] Ferdinando Fioretto, William Yeoh, and Enrico Pontelli. Multi-variable agents decomposition for dcops to exploit multi-level parallelism. In Gerhard Weiss, Pinar Yolum, Rafael H. Bordini, and Edith Elkind, editors, *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015*, pages 1823–1824. ACM, 2015.

[18] Eugene C Freuder, Marius Minca, and Richard J Wallace. Privacy/efficiency tradeoffs in distributed meeting scheduling by constraint-based agents. In *IJCAI-01 Workshop on Distributed Constraint Reasoning, August 4, 2001, Seattle, USA*, pages 63–72, 2001.

[19] Amir Gershman, Alon Grubshtein, Amnon Meisels, Lior Rokach, and Roie Zivan. Scheduling meetings by agents. In *10th International Conference . on the Practice and Theory Theory of Automated Timetabling (PATAT 2008), August 18-22, 2008, Montreal, Canada*, 2008.

[20] Amir Gershman, Amnon Meisels, and Roie Zivan. Asynchronous forward bounding for distributed cops. *Journal of Artificial Intelligence Research (JAIR)*, 34(1):61–88, 2009.

[21] Rachel Greenstadt, Barbara Grosz, and Michael D Smith. SSDPOP: improving the privacy of DCOP with secret sharing. In Edmund H. Durfee, Makoto Yokoo, Michael N. Huhns, and Onn Shehory, editors, *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, page 171. IFAAMAS, 2007.

[22] Rachel Greenstadt, Jonathan P Pearce, and Milind Tambe. Analysis of privacy loss in distributed constraint optimization. In *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, July 16-20, 2006, Boston, Massachusetts, USA*, pages 647–653. AAAI Press, 2006.

[23] Tal Grinshpoun. When you say (DCOP) privacy, what do you mean? - categorization of DCOP privacy and insights on internal constraint privacy. In Joaquim Filipe and Ana L. N. Fred, editors, *ICAART 2012 - Proceedings of the 4th International Conference on Agents and Artificial Intelligence, Volume 1 - Artificial Intelligence, Vilamoura, Algarve, Portugal, 6-8 February, 2012*, pages 380–386. SciTePress, 2012.

[24] Tal Grinshpoun and Tamir Tassa. A privacy-preserving algorithm for distributed constraint optimization. In Ana L. C. Bazzan, Michael N. Huhns, Alessio Lomuscio, and Paul Scerri, editors, *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 909–916. International Foundation for Autonomous Agents and Multiagent Systems/ACM, 2014.

[25] Tal Grinshpoun and Tamir Tassa. P-syncbb: A privacy preserving branch and bound DCOP algorithm. *J. Artif. Intell. Res. (JAIR)*, 57:621–660, 2016.

[26] Y. Hamadi, C. Bessiere, and J. Quinqueton. Backtracking in distributed Constraint Networks. In *Proc. of 13th European Conf. on Artificial Intelligence (ECAI)*, pages 219–223, Brighton, UK, 1998.

[27] Katsutoshi Hirayama and Makoto Yokoo. Distributed partial constraint satisfaction problem. In Gert Smolka, editor, *Principles and Practice of Constraint Programming - CP97, Third International Conference, Linz, Austria, October 29 - November 1, 1997, Proceedings*, volume 1330 of *Lecture Notes in Computer Science*, pages 222–236. Springer, 1997.

[28] Katsutoshi Hirayama and Makoto Yokoo. The distributed breakout algorithms. *Artificial Intelligence*, 161(1-2):89–115, 2005.

[29] Martin Hirt, Ueli Maurer, and Bartosz Przydatek. Efficient secure multi-party computation. In Tatsuaki Okamoto, editor, *Advances in Cryptology - ASIACRYPT 2000, 6th International Conference on the Theory and Application of Cryptology and Information Security, Kyoto, Japan, December 3-7, 2000, Proceedings*, Lecture Notes in Computer Science, pages 143–161. Springer, 2000.

[30] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.

[31] Sven Koenig, David Furcy, and Colin Bauer. Heuristic search-based replanning. In Malik Ghallab, Joachim Hertzberg, and Paolo Traverso, editors, *Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems, April 23-27, 2002, Toulouse, France*, pages 294–301. AAAI, 2002.

[32] Yann Krupa and Laurent Vercouter. Handling privacy as contextual integrity in decentralized virtual communities: The privacias framework. *Web Intelligence and Agent Systems: An International Journal*, 10(1):105–116, 2012.

[33] Thomas Léauté and Boi Faltings. Protecting privacy through distributed computation in multi-agent decision making. *Journal Artificial Intelligence Research (JAIR)*, 47:649–695, 2013.

[34] Michael L Littman, Stephen M Majercik, and Toniann Pitassi. Stochastic boolean satisfiability. *Journal of Automated Reasoning*, 27(3):251–296, 2001.

[35] Rajiv T Maheswaran, Jonathan P Pearce, Emma Bowring, Pradeep Varakantham, and Milind Tambe. Privacy loss in distributed constraint reasoning: A quantitative framework for analysis and its applications. *Autonomous Agents and Multi-Agent Systems*, 13(1):27–60, 2006.

[36] Rajiv T Maheswaran, Jonathan P Pearce, Pradeep Varakantham, Emma Bowring, and Milind Tambe. Valuations of possible states (vps): a quantitative framework for analysis of privacy loss among collaborative personal assistant agents. In *4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005), July 25-29, 2005, Utrecht, The Netherlands*, pages 1030–1037. ACM, 2005.

[37] Rajiv T Maheswaran, Milind Tambe, Emma Bowring, Jonathan P Pearce, and Pradeep Varakantham. Taking DCOP to the real world: Efficient complete solutions for distributed multi-event scheduling. In *3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004), 19-23 August 2004, New York, NY, USA*, pages 310–317. IEEE Computer Society, 2004.

[38] Roger Mailler. Improving the privacy of the asynchronous partial overlay protocol. *Multiagent and Grid Systems*, 8(2):165–186, 2012.

[39] Roger Mailler and Victor R Lesser. Asynchronous partial overlay: A new algorithm for solving distributed constraint satisfaction problems. *Journal of Artificial Intelligence Research*, 25:529–576, 2006.

[40] René Mandiau, Julien Vion, Sylvain Piechowiak, and Pierre Monier. Multi-variable distributed backtracking with sessions. *Applied intelligence*, 41(3):736–758, 2014.

[41] Toshihiro Matsui, Marius Silaghi, Tenda Okimoto, Katsutoshi Hirayama, Makoto Yokoo, and Hiroshi Matsuo. Leximin asymmetric multiple objective DCOP on factor graph. In Qingliang Chen, Paolo Torroni, Serena Villata, Jane Yung-jen Hsu, and Andrea Omicini, editors, *PRIMA 2015: Principles and Practice of Multi-Agent Systems - 18th International Conference, Bertinoro, Italy, October 26-30, 2015, Proceedings*, volume 9387, pages 134–151. Springer, 2015.

[42] Pragnesh Jay Modi, Wei-Min Shen, Milind Tambe, and Makoto Yokoo. Adopt: Asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence*, 161(1-2):149–180, 2005.

[43] George E Monahan. State of the art—a survey of partially observable markov decision processes: theory, models, and algorithms. *Management Science*, 28(1):1–16, 1982.

[44] Pierre Monier, Arnaud Doniec, Sylvain Piechowiak, and René Mandiau. Comparison of DCSP algorithms: A case study for multi-agent exploration. In Yves Demazeau, Michal Pechoucek, Juan M. Corchado, and Javier Bajo Pérez, editors, *Advances on Practical Applications of Agents and Multiagent Systems - 9th International Conference on Practical Applications of Agents and Multiagent Systems, PAAMS 2011, Salamanca, Spain, 6-8 April 2011*, volume 88 of *Advances in Intelligent and Soft Computing*, pages 231–236. Sprionger, 2011.

[45] Arnon Netzer and Amnon Meisels. SOCIAL DCOP - social choice in distributed constraints optimization. In Frances M. T. Brazier, Kees Nieuwenhuis, Gregor Pavlin, Martijn Warnier, and Costin Badica, editors, *Intelligent Distributed Computing V*, volume 382 of *Studies in Computational Intelligence*, pages 35–47. Springer, 2011.

[46] Adrian Petcu and Boi Faltings. A scalable method for multiagent constraint optimization. In Leslie Pack Kaelbling and Alessandro Saffiotti, editors, *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30 - August 5, 2005*, pages 266–271. Professional Book Center, 2005.

[47] Ferdinand Piette, Costin Caval, Amal El Fallah Seghrouchni, Patrick Taillibert, and Cédric Dinont. A multi-agent system for resource privacy: Deployment of ambient applications in smart environments. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1445–1446. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

[48] Stuart Russell and Peter Norvig. *Artificial Intelligence: A modern approach*. Prentice Hall, 1995.

[49] Pierre Rust, Gauthier Picard, and Fano Ramparany. Using message-passing DCOP algorithms to solve energy-efficient smart environment configuration problems. In Subbarao Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 468–474. IJCAI/AAAI Press, 2016.

[50] Julien Savaux, Julien Vion, Sylvain Piechowiak, René Mandiau, Toshihiro Matsui, Katsutoshi Hirayama, Makoto Yokoo, Shakre Elmane, and Marius Silaghi. Discsps with privacy recast as planning problems for self-interested agents. In *Web Intelligence (WI), 2016 IEEE/WIC/ACM International Conference on*, pages 359–366. IEEE, 2016.

[51] Julien Savaux, Julien Vion, Sylvain Piechowiak, René Mandiau, Toshihiro Matsui, Katsutoshi Hirayama, Makoto Yokoo, Shakre Elmane, and Marius Silaghi. Privacité dans les discsp pour agents utilitaires. In *2016 JFSMA Vingt-quatrième Journées Francophones sur les Systèmes Multi-Agents, JFSMA 2016, Rouen, France, October 5-7, 2016*, pages 129–138, 2016.

[52] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

[53] Marius C Silaghi, Djamila Sam-Haroud, and Boi Faltings. Distributed asynchronous search with private constraints. In Jeffrey S. Rosenschein Carles Sierra, Maria Gini, editor, *AGENTS 2000 Proceedings of the fourth international conference on Autonomous agents, Barcelona, Spain, June 3-7, 2000*, pages 177–178. ACM, 2000.

[54] Marius-Calin Silaghi. Meeting scheduling guaranteeing n/2-privacy and resistant to statistical analysis (applicable to any DisCSP). In *Proceedings of Web Intelligence*, pages 711–715. IEEE, 2004.

[55] Marius-Calin Silaghi and Boi Faltings. A comparison of distributed constraint satisfaction techniques with respect to privacy. In *Proceedings of the 3rd workshop on distributed constraints reasoning (AAMAS - DCR-02). Bologna, Italy*, 2002.

[56] Marius-Calin Silaghi and Makoto Yokoo. Adopt-ing: unifying asynchronous distributed optimization with asynchronous backtracking. *Autonomous Agents and Multi-Agent Systems*, 19(2):89–123, 2009.

[57] Tamir Tassa, Roie Zivan, and Tal Grinshpoun. Max-sum goes private. In Qiang Yang and Michael Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 425–431. AAAI Press, 2015.

[58] Richard J Wallace and Eugene C Freuder. Constraint-based multi-agent meeting scheduling: Effects of agent heterogeneity on performance and privacy loss. In *Proceedings of the 3rd workshop on distributed constraints reasoning (AAMAS - DCR-02). Bologna, Italy*. Citeseer, 2002.

[59] Richard J Wallace and Eugene C Freuder. Constraint-based reasoning and privacy/efficiency tradeoffs in multi-agent problem solving. *Artificial Intelligence*, 161(1-2):209–227, 2005.

[60] Toby Walsh. Stochastic constraint programming. In *Proceedings of the 15th European Conference on Artificial Intelligence*, pages 111–115. IOS Press, 2002.

[61] William Yeoh, Ariel Felner, and Sven Koenig. Bnb-adopt: An asynchronous branch-and-bound DCOP algorithm. *J. Artif. Intell. Res. (JAIR)*, 38:85–133, 2010.

[62] Makoto Yokoo, Edmund H Durfee, Toru Ishida, and Kazuhiro Kuwabara. The distributed constraint satisfaction problem: Formalization and algorithms. *IEEE Transactions on knowledge and data engineering*, 10(5):673–685, 1998.

[63] Makoto Yokoo and Katsutoshi Hirayama. Distributed breakout algorithm for solving distributed constraint satisfaction problems. In Mario Tokoro, editor, *Proceedings of the Second International Conference on Multi-Agent Systems, December 9-13, 1996, Kyoto, Japan*, pages 401–408. AAAI Press, 1996.

[64] Makoto Yokoo, Toru Ishida, Edmund H Durfee, and Kazuhiro Kuwabara. Distributed constraint satisfaction for formalizing distributed problem solving. In *Proceedings of the 12th IEEE International Conference on Distributed Computing Systems, June 9-12, 1992, Yokohoma, Japan*, pages 614–621. IEEE, 1992.

[65] Hédia Zardi, Lotfi Ben Romdhane, and Zahia Guessoum. A multi-agent homophily-based approach for community detection in social networks. In *26th IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2014, Limassol, Cyprus, November 10-12, 2014*, pages 501–505, 2014.

[66] Hédia Zardi, Lotfi Ben Romdhane, and Zahia Guessoum. Efficiently mining community structures in weighted social networks. *International Journal of Data Mining, Modelling and Management IJDMMM*, 8(1):32–61, 2016.

[67] Weixiong Zhang, Guandong Wang, and Lars Wittenburg. Distributed stochastic search for constraint satisfaction and optimization: Parallelism, phase transitions and performance. In *Proceedings of AAAI Workshop on Probabilistic Approaches in Search,, July 28 - Aug 1, 2002, Alberta, Canada*, 2002.

[68] Roie Zivan and Amnon Meisels. Synchronous vs asynchronous search on discsps. In *Proceedings of the First European Workshop on Multi-Agent Systems (EUMAS), December 18-19, Oxford, England*, 2003.